

جامعة نيويورك أبوظبي

 NYU | ABU DHABI



Lie O'clock: Experimental Evidence on Intertemporal Lying Preferences

Hande Erkut and Georgia Michailidou

Working Paper # 0076

April 2022

Division of Social Science Working Paper Series

New York University Abu Dhabi, Saadiyat Island P.O Box 129188, Abu Dhabi, UAE

<http://nyuad.nyu.edu/en/academics/academic-divisions/social-science.html>

Lie O'clock: Experimental Evidence on Intertemporal Lying Preferences

Hande Erkut

WZB Berlin Social Science Center
hande.erkut@wzb.eu

Georgia Michailidou

NYU Abu Dhabi
georgina@nyu.edu

November 2021

Abstract

In lying utility models, benefits and costs typically occur presently and simultaneously. However, lying and its products often develop asynchronously. To evaluate how these asynchronies affect the psychological costs of lying, we develop an experiment in which lying decisions occur presently, while externalities (external costs) and observability (internal costs) occur in future temporal brackets. To assess if lying costs or social preferences drive our findings, we compare against a baseline in which lying opportunities become simple distributive choices. We report significant behavioral differences when outcomes occur as products of lying rather than distributive choices which suggests that lying, per se, begets distinct psychological costs. Further, the results from exponential and quasi-hyperbolic discounting estimations suggest that temporally distancing antisocial decision-making from its consequences dilutes the associated psychological costs. External psychological costs caused by lying are discounted less and are subject to milder present-bias compared to those produced by distributive choices while manipulating internal psychological costs via observability attenuates discounting and present-bias in both cases.

1. Introduction

It feels bad when we lie. Sometimes because our lies cause harm, others because we do not like to be seen as liars. By now, there is an important body of experimental evidence showing that a non-negligible proportion of individuals will not lie for money, even if lying is not punishable or hurtful to others. A lot of research has been dedicated to understanding the roots of this reluctance to lie, which seems inelastic to

monetary incentives but sensitive to psychological controls (Kajackaite and Gneezy 2017).¹ The latest consensus suggests that we not only entertain a preference for honesty per se, but also, a preference for appearing honest. That is, the lying costs include reputational concerns (Gneezy, Kajackaite and Sobel, 2018) (Abeler, Nosenzo and Raymond, 2019). Albeit complex, the psychological costs of lying might be the key for controlling the presence and magnitude of deceptive practices in the economy. This paper is about these lying costs.

Using an economic experiment, we set out to find out *when* and *how* lying costs are experienced. First, we ask *when* the psychological costs are endured because lying is not always instantaneous. It is not often the case that individuals come upon opportunities to lie about an outcome that just occurred and receive the costs and benefits on the spot. Lying can take many asynchronous forms. For example, there can be concealing lies that individuals use to conceal past actions (consider a guilty suspect lying about their whereabouts three weeks ago, during the time of the crime). Or, there can be anticipatory lies that individuals use on a present time while expecting the benefits in the future (for example, lying about your CV qualifications during a review hoping to get the company promotion next month). If we think of lying decomposed in four-time elements (a) the moment the truth becomes, (b) the moment the lie is expressed, (c) the moment its benefits are received, and (d) the moment the costs are suffered, then many more time asynchronies arise. However, the economics literature on lying has only focused on instantaneous lies². In this paper, we take the next step towards studying some of the intertemporal properties of lying.

In particular, we examine whether temporally distancing lying from its negative consequences affects the intensity with which the associated psychological costs are experienced. To fix ideas consider Jack, a health insurance agent, who lies about the insurance package he is selling. He gets money now and some years later, the clients find out they are not covered for many important health services. Does Jack feel any less bad about his lie because the externalities for the clients come later? We know from economics, psychology, and the neurosciences that temporal distances change the way we think of events. From motivated memory (Saucet and Villeval 2018) to delay discounting (Odum 2013) (Bickel, et al. 2014), we make choices differently when a temporal component enters our decision functions. Yet, none of the

¹ The reluctance to lie has been modeled in different ways. Some postulate that honesty is a strong moral norm that triggers strong feelings of shame and guilt to those who lie (Bicchieri, Dimant and Sonderegger 2020). Others conceptualize the reluctance to lie as an explicit preference for honesty (Kartik, Tercieux and Holden 2014) or to lying costs (Abeler, Becker and Falk, 2014).

² An exception being (Andersen, et al. 2018) which introduces reflection time by distancing observing outcomes from reporting about them.

existing temporal models can capture the case of Jack fully. Let us consider the standard models of discounting. In those models, agents experience their utility the moment its components materialize. That is, in a choice between \$10 now and \$12 in a week, agents compare the utility of consuming \$10 now to the utility of consuming \$12 in a week. They might discount the \$12 by assigning it a lower present value, they might even factor in anticipation, but they do not explicitly formulate present day feelings about their future consumption. Said differently, in these discounting models, the utility from consumption emerges when the consumption materializes – not before. If we apply the same reasoning to Jack's context, and assume that the externalities of his lie cause him some disutility, this disutility should be experienced the moment the externalities materialize. Therefore, Jack should not be feeling bad the moment he lies. Rather, he should feel bad years later, when his clients get hurt. But, it is easy to think of alternative emotional time paths. Maybe Jack is not calculating a present value for his future disutility to evaluate whether it is worth lying right now. Maybe he is feeling bad about lying –right now- despite the fact that his lies have not yet harmed anyone. If so, how should we go about theorizing Jack's decision-making timeline?

Due to the complexity of lying costs, and intrinsic motivations in general, it is difficult to formulate assumptions about the timing of psychological utility obtainment. In this paper, we propose an estimation of this timing by means of relative comparisons. First, guided by the existing literature, we construct an experimental set up in which externalities and observability, two of the most prominent lying cost triggers, can emerge. Then, we introduce a temporal component that distances the materialization of externalities and observability from the moment of lying expression. Subsequently, using exponential and quasi-hyperbolic discounting models, we compare the intensity with which these psychological costs are experienced and whether individuals discount them differently. The results from this exercise help us understand if the mere distancing of externalities and observability has any effect on current lying decision-making processes. However, this exercise does not help us understand fully *how* these costs are experienced. Assume for example that the disutility from lying occurs the moment the lie is expressed and not when the externalities or observability occurs. This would manifest with a rather low discounting rate and present bias. However, such a result cannot be attributed to the intertemporal properties of lying – alone- . One could argue that distancing any antisocial choice from equivalent consequences might yield the same result, i.e., it is not the lying aspect of the choice that draws the emotions to the present, but rather the negligible discounting of the later-in-time antisocial outcomes.

To evaluate whether this is the case, one needs to see how equivalent (anti)social preferences develop intertemporally and make the relevant comparisons. In particular, if in a similar set up where outcomes emerge as products of choices, not lying, individuals discount these outcomes identically, one would be right to argue that, indeed, the timing of the outcomes drives utility obtainment. Conversely, if the same outcomes were subject to different discount rates, then we would have evidence that lying per se is responsible for distinct utility dynamics. A review of the findings on intertemporal social preferences could provide a guide for this. However, the literature on social preferences still offers very few insights on the intertemporal development of these preferences although this is changing with a few recent papers. Despite the plethora of empirical regularities collected on instantaneous cooperation, altruism, or trust and reciprocity, we are still short in documenting the intertemporal component of social preferences in similar detail. Questions like whether we care about future inequalities as we care about present ones, or whether reciprocation scales down with time are still a long way before conclusively answered.

To help us establish whether the psychological costs of lying maintain distinct temporal properties from equivalent antisocial actions we extend our experimental set up to include a *ceteris paribus* decision environment in which lying opportunities are removed. We have additional reasons to want to understand whether lying is experienced just as any other antisocial action or whether it imprints a distinct psychological trail. If extant, this distinction can help us contextualize relevant policy implications. If people feel bad because of the consequences their actions/choices carry, then generalized policies that leverage those consequences should be proposed/designed. However, if people feel bad about the way these consequences come to be, i.e., they are sensitive to whether such consequences are the products of lying in particular, then more specialized policies that embed the moral extensions of lying need to be considered.

In our experimental set up, there is a sender and a passive receiver. The sender observes an outcome and decides whether to send an honest or not honest message about it to the receiver. Honest messages pay an equal amount to both players while dishonest messages reduce receiver's but increase sender's pay. Though their own payment always occurs in present time, senders make consecutive decisions about sending honest or not honest messages while the receiver's payment moves to further dates. In an observability treatment, senders and receivers' payment always occurs in present time but receivers eventually find out whether senders were honest. Senders make consecutive honest or not honest decisions while the revealing of their (dis)honesty moves to further dates. To estimate precisely the extent to which any observed discounting is endemic to lying or driven by distributive preferences, we develop

two additional treatments in which lying decisions become pure allocation decisions following the same time dynamics. Our results demonstrate a significant contrast across the lying and allocation treatments. Psychological costs related to certain outcomes are significantly higher when these outcomes occur as products of lying rather than allocation decisions. Further, the lower psychological costs derived from antisocial allocation decisions are subject to significantly higher discounting and present bias than those derived from lying. The psychological costs of lying are subject to very low discounting and present bias. Individuals discount costs derived from future externalities very little while they do not discount at all reputational costs produced by future observability. Put simply, our results suggest that lying carries distinct psychological cost properties, statically and intertemporally. When people make honest/dishonest decisions, they experience the associated negative feelings the moment lying is expressed, irrespectively of when the negative products of lying occur.

These findings are important for both academic and policy design reasons. Primarily, they call for further investigation on intertemporal social preferences. Contrarily to standard economic intuition we show that, time delays are not discounted equivalently for equivalent monetary incentives. This in turn suggest that we need a more refined way of theorizing intertemporal preferences. One that accounts not only the betas and deltas of monetary outcomes but also how these outcomes come to be. More importantly, we need ways of expressing how different means of outcomes trigger distinct discounting processes and we need to document empirical regularities on these distinct discounting processes more rigorously. Policy related, our findings suggest that individuals are sensitive to contextual parameters. Potentially, choice architectures that highlight the antisocial aspect of certain actions might not enough to curb such actions. However, if such actions can be framed in the context of lying then stronger psychological costs can be triggered with the added benefit of them being time invariant.

The manuscript proceeds as follows; in section 2 we review the literatures on lying, intertemporal preferences, intertemporal social preferences, time perceptions, and methodologies pertaining to the one we apply here. Section 3 introduces the game and section 4 an ad hoc model for the game we consider. We describe extensively the experimental design on section 5 and provide the results in section 6. Section 7 concludes the manuscript.

2. Contribution and Relate Literature

Our paper extends contributions to both, academic literatures and policy designing. First, we quantify the psychological costs of lying, a pivotal parameter in ethical decision-making contexts. We structurally estimate these costs statically and intertemporally using data collected from a carefully designed economic experiment. Then, given its extension to literatures on (anti)social preferences and intertemporal preferences, our paper illustrates the need for further and more rigorous research on intertemporal social preferences. Finally, policy related, our findings suggest that individuals are sensitive to contextual parameters. Potentially, choice architectures that highlight the antisocial aspect of certain actions might not enough to curb such actions. However, if such actions can be framed in the context of lying then stronger psychological costs can be triggered with the added benefit of them being time invariant.

Turning to the literature review, In the lab lying behavior has been studied extensively using the die-roll paradigm introduced by (Fischbacher and Föllmi-Heusi 2013). In their design, subjects privately observe the outcome of a die-roll and are asked to make a report that will determine individual monetary payoffs. Fischbacher and Follmi-Heusi (2013), as well as (Gächter and Schulz 2016) and Gneezy et al. (2018) find that people lie but not maximally. Lying behavior has also been investigated under different types of strategic interactions, such as deception games and coordination games. In deception games, subjects strategically choose to lie taking into consideration not only their own gains but also the possibility to favor or damage others with their choice (Gneezy, 2005; Erat and Gneezy, 2012; Gneezy et al., 2013). In coordination games, people strategically lie more if the monetary benefits of lying are equally shared with others i.e. there is a payoff commonality (Conrads et al., 2014; Weisel and Shalvi, 2015; Barr and Michailidou, 2017) or when they can communicate (Kocher, Schudy and Spantig 2016). In coordination games where the benefits of lying are not equally shared, Lauer and Untertrifaller, 2019 find that one third of the subjects engage in strategic lying if and only if one or more group members lie.

The lying game is a hybrid of various lying games proposed in the literature. It has the general structure of sender-receiver games as those in Gneezy 2005 and Gneezy et al. 2013 but receivers have no power of choice, making senders dictators with lying opportunities. It uses a binary distribution of true states as Abeler et al. 2019 and employs a skewed probability distribution over true states as (Munoz and Makovi 2021), while it borrows aesthetic elements from (Alempaki, Burdea and Read 2021).

This paper makes a contribution to the long-standing literature investigating intertemporal preferences (see Frederick et al. (2002) and Cohen et al. (2017) for extensive reviews of the literature), as well as the small and recent literature on intertemporal and social preferences . The literature on intertemporal and

social preferences provides mixed evidence: Bremen (2011) and, Andreoni and Serra-Garcia (2019) show that people give more to charities when the giving is delayed as opposed to taking place today. This evidence is in line with the choice-dated utility approach. On the other hand, Kolle and Werner (2021) does not find evidence for time-inconsistent giving, and Kovarik (2009) and Dreber et al. (2016) find evidence for lower rates of giving when the giving is delayed in a dictator game context. Recently, Chopra et al. (2021) show that results from their intertemporal charitable giving experiment can only be explained by a utility function that has both a consequence-dated and a choice-dated component.

3. The Baseline Experimental Set Up

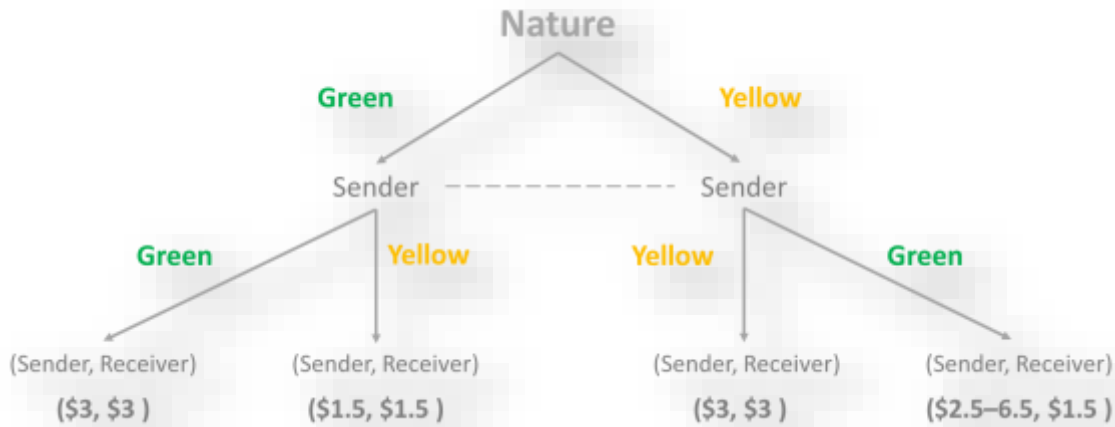
To describe the baseline experimental set up, let us introduce Players = {Sender, Receiver} and let Sender $s = \{1, 2, 3, \dots, N\}$ and Receiver $r = \{1, 2, 3, \dots, N\}$. Players are paired for the purposes of the game, in a random fashion, and have no information about each other's persons. There is a set of states $K = \{G, Y\}$ where G stands for green and Y for yellow. Sender observes a signal $O_S \in K$ drawn and sent by Nature. Nature draws G with probability p and Y with $1-p$. Green signals are much less frequent and while players know that, they do not know the exact probabilities³. Upon observing O_S , sender is instructed to send an electronic Message $M_S \in K$ to the receiver which, although requested so, needs not to match O_S . Thus, the strategy space K^2 of the sender has four elements (O_S, M_S) that are all the possible combinations of signals and messages that can be received and sent. Coming to the incentives of the game, senders and receivers gain \$3 if $M_S = O_S$. Senders decide whether $M_S \neq O_S$, for amounts increasing progressively from \$2.5 to \$6.5, while receivers in this case gain \$1.5.

In this baseline version, only the sender knows, and knows she is the only one who knows, the full payoff structure. The receiver is agnostic about the payoff structure and of the lying opportunities of the sender. This means the honesty of the sender is not verifiable by the receiver; however, receivers knowledge of how payoffs come to be, as well as the observability of honesty is a varying feature across our treatments.

Describing this game differently, the senders gain a high amount if they send a green message and a low amount if they send a yellow message irrespectively of the true state of nature, while receivers gain a high amount if the message they receive is an honest message and a low amount otherwise, irrespectively of

³ We know from Gneezy et al. 2018 and from Abeler et al. 2019 that changing the probability of drawing high-payoff states affects lying. Here, we choose to have Sender being agnostic about the exact probability in order to avoid confounding any such effects with the factors we are interested. Senders might form beliefs about these probabilities, but there is no reason to suspect that the distribution of beliefs will be different across treatments therefore we expect the disperse of these beliefs to be constant. In the experiment, a Green state was drawn with probability 3%.

the content of the message. As follows, Nature drawing green would not involve any ethical dilemma from the point of a rational sender, and would imply the highest collective benefit for the pair. Contrarily, Nature drawing yellow puts the sender in a dilemma of either sending an honest message and forgoing a monetary gain, or lying and causing a monetary loss to the receiver and potentially experience a damage to his self-image. Below, we present the tree graph of this game and then proceed to compose an ad hoc model that describes senders' utility.



Graph 1: Baseline Decision Tree

4. The model

Senders have a monetary incentive to send a Green message, irrespectively of the state Nature drew⁴. However, when Green message is a lie, receivers suffer monetarily. Further, in some variations, senders' truthfulness becomes known to the receiver. Thus, this game involves factors that might affect lying choices in different ways. These factors are the externalities lying carries, and the potential observability of lying. In a series of treatments, we vary the timing of these factors' occurrence and we estimate both, the sensitivity of lying to each of these factors, and their discounting rates. Below we use an ad hoc model to help us consider the various possible contexts and timings our baseline set up could expand upon.

Let $U_S: K^2 \rightarrow \mathbb{R}$ be the utility function of the sender. We assume this is given by:

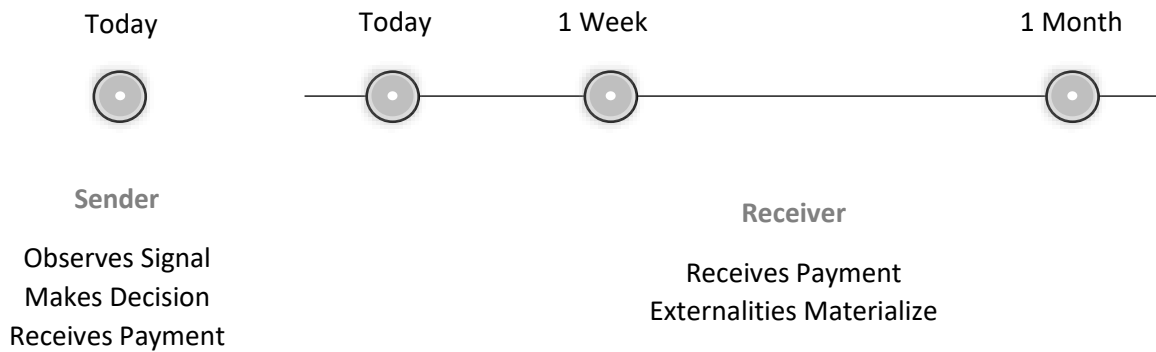
⁴ There are two states for which Green is not the highest paying option. These serve only as sanity checks and have no strategic interpretation in our set up.

$$U_s[I, L, E, O] = U(I) - D_t C(L, E, O)$$

Let us call the above utility function UF1 and define $I \in (\$1.5, \$6.5)$ as the parameter that captures the monetary *Incentives* of the sender associated with his choices. Notice that, in our set up, senders' payment only occurs in present time therefore monetary incentives are not discounted. We give senders a list of present-time monetary choices and ask for which level of payment they are willing to switch from making a pro-social to making an anti-social choice. We use these switching points to estimate idiosyncratic *Cost* parameter $C \geq 0$, and we document how C varies when we consider the contexts and timings given by the following parameters:

- $L \in (0, 1)$, defines a lying context. We evaluate the extent to which costs are endemic to lying or driven by distributive preferences by comparing choices about honesty/dishonesty to corresponding choices about equal/unequal allocations. Therefore, we use context $L = 0$, to benchmark how psychologically costly inequalities are, and context $L = 1$ to measure how costly these inequalities become when produced by lying.
- E is the externalities parameter. Externalities remain fixed for all contexts and timings we consider. An antisocial choice, either as the product of lying or mere allocation, harms the receiver by \$1.5. What we vary with respect to externalities is their timing. We push the emergence of externalities in future dates. Essentially, we are measuring how psychologically costly it is to hurt a receiver by \$1.5, whether distancing the damage in further dates affects the associated psychological cost and consequently the willingness to act antisocially.
- $O \in (0, 1)$ is the parameter that captures the context of observability. Motivated by the related literature, we evaluate whether the psychological cost of causing negative externalities differs if choices are observable $O = 1$ or not observable $O = 0$. Further, we estimate whether choices differ if observability occurs in future dates.

Lastly, D_t is the parameter that captures the discounting of the psychological costs depending on the different contexts we consider with our experimental treatments. It allows that different combinations between lying contexts, externalities, and observability be subject to different discounting. What is constant among all our timing manipulations is the time intervals between choices and later outcomes. In the timeline below, we show how all decisions are made in present times, but concern outcomes that occur in present time, in one week, and in one month (30 days). In the following section, we describe extensively the experimental design and we customize the timeline to each experimental treatment.



5. Experimental Design

Now, we describe the method we will use to estimate each of the parameters presented in the model and their respective discounting rates. We employ a 2x2 experimental design across which we vary parameters L lying, and O observability. Subjects participate only in one of the four treatments. The timing of consequences varies within each treatment. Thus, time preferences are elicited within subjects while behavioral reactions to parameters L and O are elicited between subjects. We will call treatments with observability as Observability treatments and those without observability as Externality treatments. Similarly, the treatments with lying opportunities as Lying treatments, and those without, Choosing treatments. The table below gives the treatment synopsis. Videos with experimental instructions and choice interfaces for all treatments can be accessed “here”.

	Lying	Choosing
Externalities	$O = 0,$ $L = 1$	$O = 0,$ $L = 0$

Observability	$O = 1,$ $L = 1$	$O = 1,$ $L = 1$
---------------	---------------------	---------------------

5.1. Treatment Externalities + Lying

Subjects visit the online platform and in a common information screen, they read the following:

Instructions

We are now ready to proceed with the instructions. Read these instructions carefully as there is an understanding question on the next screen. **If you do not answer this question correctly, your participation in the study will end and you will not receive payment.**

In this study, you will be matched with another participant to take part in a "color game".

One participant will be **Member A**, the other **Member B**.

In the game, a randomization algorithm will draw a color, either **GREEN** or **YELLOW**. It is less likely that the algorithm will pick **GREEN** than **YELLOW**.

The color will be seen only by Member A.

Member A's task is to send Member B a message indicating the color drawn by the algorithm. **The rule is that Member A is honest so that Member B receives an honest message.**

The color drawn by the algorithm and the honesty of the message are important because they determine the bonus payments for Member A and Member B.

Member B will not make any choices and will not know the precise choices available to Member A. However, Member B always earns more if Member A sends an honest message. **Member B will not be informed if the message sent by Member A was honest or not honest.**

In the next screen, you will find out if the randomization algorithm assigned you to the role of Member A or Member B. You have equal chances of being assigned to each role.

After the role assignment, Member A observes the color and, if yellow, he is invited to make 45 choices⁵. For each choice, choosing left sends an honest message, while choosing right sends a not honest message. There are three sets of 9 left/right choices, each paying Member B today, in a week, or in a month.



Notice that, for the two future choice sets, the payment of Member B moves to the future for both left and right choices of Member A. Had we kept the returns from honesty payable to Member B today but returns from dishonesty payable in the future, we would have failed to account for any present bias Member A applies to the monetary incentives of Member B. This would have inflated the present value

⁵ If the color observed is Green, Member As make only one choice between an honest message that pays both players \$3, or a not honest message that pays both players \$1.5. By construction, only 3% of Member As observed Green and all of them, in all treatments, sent an honest message. Member As' behavior when observing Green serves no purpose in our analysis thus we will not be making further reference to them in the rest of the manuscript.

of the honest choice, making the monetary aspect of left/right choices unequal between the today and the future choice sets.

Once they respond to the three time varying choice sets, Member As respond to two additional choice sets in which they make choices between a certain option (left choice) of \$3 to both players or a risky option (right choice) which pays a fixed lower amount to Player B and a higher amount to themselves with varying probabilities. The lower and higher amounts presented in the risky column are tailored to the switching points Member As exhibited in the time choice sets. That means, Member As who switched from yellow to green messages earlier on in the time choice sets, view risky options with smaller gap between own and Member B payments. This sensitivity in the risk elicitation helps us elicit risk preferences more accurately within the population of Member As. In turn, the reason we elicit risk preferences is to allow us to map utility curvatures for the structural, maximum likelihood estimations we will perform later on. As shown in Andersen et al. (2008), estimations that do not consider utility curvatures might be subject to overestimation of discounting rates.

Send the honest message YELLOW which pays:	Send the not honest message GREEN which pays:	Send the honest message YELLOW which pays:	Send the not honest message GREEN which pays:
\$3.00 to me and \$3.00 to Member B	\$4.00 to me with 100% probability or \$6.00 with 2% probability and \$1.50 to Member B	\$3.00 to me and \$3.00 to Member B	\$3.00 to me with 100% probability or \$7.00 with 1% probability and \$1.50 to Member B
\$3.00 to me and \$3.00 to Member B	\$4.00 to me with 87.5% probability or \$6.00 with 12.5% probability and \$1.50 to Member B	\$3.00 to me and \$3.00 to Member B	\$3.00 to me with 87.5% probability or \$7.00 with 12.5% probability and \$1.50 to Member B
\$3.00 to me and \$3.00 to Member B	\$4.00 to me with 75% probability or \$6.00 with 25% probability and \$1.50 to Member B	\$3.00 to me and \$3.00 to Member B	\$3.00 to me with 75% probability or \$7.00 with 25% probability and \$1.50 to Member B
\$3.00 to me and \$3.00 to Member B	\$4.00 to me with 62.5% probability or \$6.00 with 37.5% probability and \$1.50 to Member B	\$3.00 to me and \$3.00 to Member B	\$3.00 to me with 62.5% probability or \$7.00 with 37.5% probability and \$1.50 to Member B
\$3.00 to me and \$3.00 to Member B	\$4.00 to me with 50% probability or \$6.00 with 50% probability and \$1.50 to Member B	\$3.00 to me and \$3.00 to Member B	\$3.00 to me with 50% probability or \$7.00 with 50% probability and \$1.50 to Member B
\$3.00 to me and \$3.00 to Member B	\$4.00 to me with 37.5% probability or \$6.00 with 62.5% probability and \$1.50 to Member B	\$3.00 to me and \$3.00 to Member B	\$3.00 to me with 37.5% probability or \$7.00 with 62.5% probability and \$1.50 to Member B
\$3.00 to me and \$3.00 to Member B	\$4.00 to me with 25% probability or \$6.00 with 75% probability and \$1.50 to Member B	\$3.00 to me and \$3.00 to Member B	\$3.00 to me with 25% probability or \$7.00 with 75% probability and \$1.50 to Member B
\$3.00 to me and \$3.00 to Member B	\$4.00 to me with 12.5% probability or \$6.00 with 87.5% probability and \$1.50 to Member B	\$3.00 to me and \$3.00 to Member B	\$3.00 to me with 12.5% probability or \$7.00 with 87.5% probability and \$1.50 to Member B
\$3.00 to me and \$3.00 to Member B	\$4.00 to me with 0% probability or \$6.00 with 100% probability and \$1.50 to Member B	\$3.00 to me and \$3.00 to Member B	\$3.00 to me with 0% probability or \$7.00 with 100% probability and \$1.50 to Member B

A feature of our design is the multiple price listing (MPL onwards) style of the choice sets. One can argue that, when applied to ethical decision making dilemmas, this feature might induce cold decisions, i.e., trigger logic and critical analysis processes that could potentially reduce the moral and put more emphasis on the quantitative component of the dilemma. Our way of maintaining the moral aspect of the decisions

is to highlight repeatedly that left-side choices are *honest* choices while right side choice are *not honest*, while also emphasizing that the rule of the color game is to send an honest message. We decided against referring to *not honest* choices as lying or dishonest choices because we conjecture this would be a very strong overtone, stretching the moral aspect more than necessary.

Another feature of our design is that the monetary incentives are constant across all treatments. This means that the extent to which efficiency, inequality aversion, or other type of social preference drive our results remains constant across treatments, therefore, any treatment differences can solely attributed to the variables we manipulate within each treatment.

We incentivize decisions by paying one randomly drawn choice from all the 45 left-right choices Member As make. As proposed by the literature (Azrieli et al. 2015, Charness et al., 2016) this “pay one” incentivization approach controls for cross-task contamination, wealth effects, hedging, and bankruptcy considerations that might have risen if all choices were payable.

Our findings from this treatment will serve as a benchmark of comparison. First, we will estimate the value of the psychological cost incurred when deciding on an action that carries negative externalities for another. Specifically, when these externalities are the product of lying and lying is not observable. Then, we will estimate the discounting of this psychological cost, i.e., we will see whether lying feels less hurtful the further away externalities occur in the future. Finally, we will estimate the present bias of the psychological costs, i.e., we will see the sensitivity with which the negative feelings from lying contract when the externalities are removed from an immediate time bracket to a future one, irrespectively of the actual timing of the future. In other words, the present bias will give us a perspective on how much stronger the psychological costs are experienced when the externalities occur today compared to any future date.

Comparing the parameters estimated from this treatment to the equivalent parameters estimated from treatment Observability+Lying, will show us how the psychological cost of lying are affected by social image considerations. That is, whether introducing transparency of lying decisions affects the intensity with which costs are incurred and whether the cost time dynamics are affected. Further, when comparing with the Externalities+Choosing treatment, we will be able to assess the extent to which the costs we document are endemic to lying, or, are driven by distributive preferences. In particular, if any comparison yields significant differences, these differences could be attributed to the one-only point of divergence between the two treatments; how outcomes come to be. In this treatment, outcomes are the products

of lying while in the Externalities+Choosing, the same outcomes occur as choices. Therefore, any treatment differences will manifest the moral attachment lying comes with in ceteris paribus decision environments.

5.2. Treatment Observability+Lying

This treatment is the same as Externalities+Lying, except from the fact that Member A's choice between an honest or not honest message will eventually be disclosed to Member B's. What we vary, is the timing of the choice disclosure. While payments for both Members are payable today, Member A makes left/right choices knowing that Member B will find out about the honesty of the message today, in a week, and in a month.

5.3. Treatment Externalities+Choosing

This treatment is identical to the Externalities+Lying treatment, however, there are no longer honest/not honest dilemmas but simple left/right choices. We have kept the description of the color game the same but now players are informed that a green signal means Member A's need to make only one choice, while a yellow signal means Member A's need to make 45 choices. The general instructions screen subjects see in this treatment is the following

Instructions

We are now ready to proceed with the instructions. Read these instructions carefully as there is an understanding question on the next screen. **If you do not answer this question correctly, your participation in the study will end and you will not receive payment**

In this study, you will be matched with another participant to take part in a "color game".

One participant will be **Member A**, the other **Member B**.

Member A will choose the bonus payments of Member A and Member B.

In the game, the randomization algorithm will draw a color, either **GREEN** or **YELLOW**. It is less likely that the algorithm will pick **GREEN** than **YELLOW**.

The color will be seen only by Member A.

The color drawn by the algorithm is important because it determines the bonus payments Member A can choose from.

Member B will not make any choices and will not know the precise choices available to Member A. **Therefore, Member B will not be informed whether Member A gave him or her the highest available bonus payment.**

In the next screen, you will find out if the randomization algorithm assigned you to the role of Member A or Member B. You have equal chances of being assigned to each role.

5.4. Treatment Observability+Choosing

This is a treatment merge between treatments Observability+Lying and Externalities+Choosing. Similarly to Externalities+Choosing there are no lying opportunities but choices – and similarly to Observability+Lying, both member's receive payments today but Member A's choice (whether it is one that helped or hurt Member B) is disclosed to Member B today, in a week, and in a month. Treatment Observability+Choosing helps us understand the source of social image concerns, an intrinsic motive that has been identified as an important determinant of lying decisions. In particular, the comparison with treatment Observability+Lying, when observability occurs today, will reveal whether social image

concerns are fueled by observability per se or observability of lying in particular. Second, the intertemporal comparisons will help us evaluate whether social image concerns follow distinct intertemporal development depending on the source of these concerns.

6. Procedures

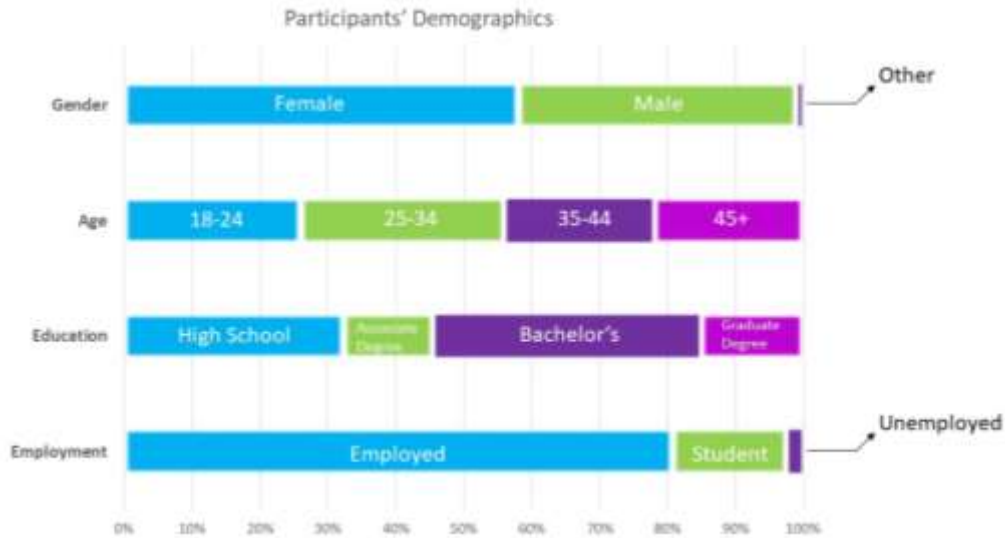
We executed all sessions of this experiment online on the 15th of April 2021. We generated the interface on Qualtrics and recruited Prolific subjects residing in Great Britain. We particularly chose Prolific because its payment system had specific advantages that could accommodate the complexities of the experiment. Specifically, the experiment required payments and individualized messages to be sent on different points in time. Prolific has a subject-individual payment system through which bonus payments can reach subjects in specified future times. Also, a messaging system which allows the experimenter to message subjects individually and anonymously at different points in time. We only recruited subjects with high Prolific approval rate, which measures the success of participation in other studies. This means that, having participated in other studies, subjects were familiar with the platform and could trust the reliability of the payment system. Although the design of our experiment did not involve decisions about sooner versus later payments (for which subjects' trust on the delivery of future payments is critical) it was still important to convey that any future-future decision they made would be honored to its entirety. Also, the fact that our design required subjects to actively participate only once restricted potential self-selection issues that arise when multiple engagement dates are required, often the case with intertemporal choice experiments.

Throughout their participation, subjects had to answer pre-announced attention checks. Any one wrong answer disqualified participation. Subjects could not log in and try to participate again either upon completion of disqualification as the platform could detect and restrict these subjects. It is highly unlikely that bots can penetrate Prolific registration system and the fact that we selected subjects who had participated in previous research together with the attention checks we implemented makes it almost impossible to have any bot decisions among our data.

1000 subjects completed successfully our experiment, approximately 250 in each treatment condition⁶. Allocation of subjects among conditions was randomized on a subject level. The order of time choice sets

⁶ Randomization on a subject level, combined with disqualification of subjects who did not pass attention checks produces slight variations in the exact number of subjects across treatments. Excluding those subjects who observed "Green" color, the break down of subjects across treatments is: Externalities+Choosing (246),

presented was also randomized on a subject level with some subjects viewing today sets first, others one-week sets first, and so on. Subjects' demographics distribution does not vary significantly between any two treatment comparisons. A graph with the basic demographics among all subjects is below.




7. Results

To generate a general view of how our results developed across treatments we plot the following graph which depicts the psychological costs associated with antisocial actions, be them choices or lies, depending on the treatment.

Observability+Choosing (264), Externalities+Lying (188), Observability+Lying (254). Out of all subjects in each treatment, half were Member As and half, passive Member Bs.

Psychological Costs in Monetary Units



In the graph above, we use the switching points from each time and treatment condition to depict the psychological cost of the associated antisocial action and we represent the 95% confidence intervals with the purple sticks. In more detail, the today bar, under the Externalities+Choosing treatment, shows that, on average, subjects switched from left (prosocial) to right (antisocial) choices between rows 5 and 6 (see time choice sets graph in the section above). That means they were willing to forgo £1.78 before switching to a state that harms Member B by £1.5 but pays them more. Thus, switching to lower rows is associated with higher amounts forgone and therefore higher psychological costs as there is no other reason prohibiting earlier switches.

7.1. Static Psychological Costs Across Treatments

7.1.1. Effect of Lying

Comparing today bars between Externalities+Choosing (£1.78) and Externalities+Lying (£2.85), we find the difference in psychological costs strongly significant ($p < 0.01$). This finding indicates that lying, alone, begets distinct psychological costs that are additive to the costs an equivalent state of incentives produces when reached via choosing rather than lying. The same result is obtained when we compare how psychologically costlier lying is when we compare the treatments involving observability. The difference between Observability+Choosing (£2.17) and Observability+Lying (£3.15) is also strongly significant

($p < 0.01$). Taken together, these two results bring strong evidence on how taxing the concept of lying is in decision making.

7.1.2. *Effect of Observability*

Turning to the effect of observability, i.e., social image concerns, first we compare Externalities+Choosing (£1.78) with Observability+Choosing (£2.17). The difference is statistically significant ($p = 0.02$) suggesting that the disclosure of Member As choices to the affected Member Bs imposed a psychological burden to Member As, making them more reluctant to act antisocially.

This social image burden goes in the same direction when lying decisions are involved, albeit milder. When comparing the psychological costs between Externalities+Lying (£2.85) and Observability+Lying (£3.15), the difference is marginally statistically significant ($p = 0.09$). This can be explained either as, indeed, stronger image concerns when individuals decide rather lie over outcomes, or as ceiling effects. That is, observability over lying outcomes might bring as strong or stronger social image concerns, however, the scale of switching rows in our experiment might be too short to capture it.

7.2. Intertemporal Psychological Costs Within Treatments

Having shown the effect of lying and observability in decision making, we now turn to evaluate how decision making develops in intertemporal contexts. First, we will study the intertemporal development by testing whether the costs, as depicted in the graph above, vary across the tree time brackets we consider. In subsequent sections, we will perform structural estimations to measure discount factors and present bias parameters.

7.2.1. Choices with Externalities in the Future

We begin to evaluate whether the psychological costs of antisocial decisions reduce when the consequences of these decisions occur in the future. In our context, we will be testing if Member As are willing to engage in antisocial actions for smaller profit when the externalities to Member Bs are born in the future. In treatment Externalities+Choosing, we compare the bar of today (£1.78) with the bar of 1 week (£1.51) and find the difference statistically significant ($p = 0.002$), however, we find no further decrease in the accepted profit when we move further in time. As also evident from the graph, moving the consequences of antisocial choices away from today, leads to an increase of antisocial choices; however, the time distance from today does not seem have a further impact on these choices.

7.2.2. Choices with Observability in the Future

We repeat the exercise for treatment Observability+Choosing and observe a similar pattern. When comparing today (£2.17) with 1 week (£2.02) we find the difference in antisocial choices significant ($p=0.051$), however, moving the observability of the antisocial choices further in the future, does not seem to impact psychological costs further as the difference in antisocial choices does not increase.

7.2.3. Lying with Externalities in the future

Turning to the treatments involving lying, time dynamics seem different. Lying choices are invariant to the timing of externalities. That is, when considering treatment Externalities+Lying, subjects do not lie in a significantly different manner when the externalities of lying occur today, in a week, or in a month, suggesting that the psychological costs associated with lying are not sensitive to the timing of consequences; rather they seem to be drawn and experienced in the moment lying occurs.

7.2.4. Lying with Observability in the future

Similarly, varying the observability of lying across time does not seem to impact lying decisions. When considering treatment Observability+Lying, psychological costs, seem to be drawn to the moment of lying and not on the moment when observability of lying occurs. This brings further evidence that, while antisocial choices in general are sensitive to the timing of consequences, lying decisions are not.

7.3. Structural Estimations.

We now turn to structural estimations to consider more formally discounting rates and present biases. These type of estimations typically require assumptions about some structural form of a utility function. In what we show below, we follow assumptions and procedures similar to (Hey and Orme 1994), (Holt and Laury 2002), (Andersen, et al. 2008). We are fixing a random utility choice model and estimating it by applying maximum likelihood estimations. To do so, we are assuming that the stochastic errors of the random choice model are normally distributed.

We perform a series of estimations where we vary the structural form of the utility function depending on whether we account for risk preferences and on whether we assume hyperbolic or quasi hyperbolic discounting. We also perform robustness checks for these estimations in which we control for corner solutions and time choice set order effects. We present these checks in the appendix.

The utility function parameters we will be estimating throughout this section are the following

$\kappa \geq 0$	The psychological cost parameter. Before, we used switching points to measured psychological costs in monetary terms. Now we estimate these costs as utility parameters. In particular, κ measures the psychological cost of choosing right instead of left in the time choice sets
$\delta > 0$	The discount factor. It will show us the rate with which κ declines, when externalities and observability move in the future.
$\alpha \geq 0$	Coefficient of relative risk aversion. It captures subjects' risk preferences and it is used to define the curvature of the utility function
$\beta \geq 0$	Present bias parameter. When assuming quasi-hyperbolic discounting, this parameter typically shows the impatience of individuals for instantaneous outcomes – in our context it shows the sensitivity of psychological costs when externalities and observability are drawn from the moment of decision to any future moment

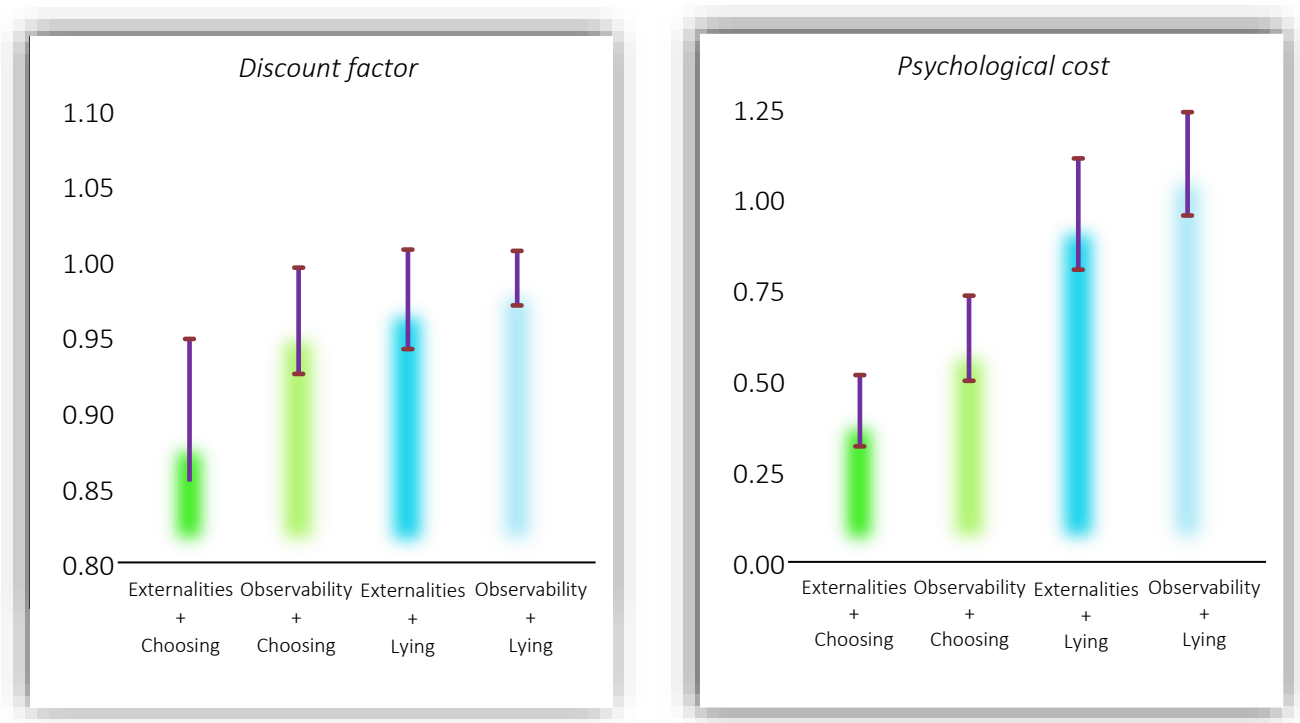
In the table below, we present a summarizing view of all the results from estimations. While we discuss graphically and analytically columns 2 and 4 in the sub-sections below, we present the rest of the columns in the appendix.

	(1) Risk Neutral Exponential Discounting	(2) CRRRA Exponential Discounting	(3) CRRRA Exponential Discounting (robustness I)	(4) CRRRA Exponential Discounting (robustness II)	(5) CRRRA Quasi- Hyperbolic Discounting
Psychological cost κ					
Observability+Choosing	0.132* (0.053)	0.200** (0.077)	0.051 (0.049)	0.106 (0.068)	0.199** (0.079)
Externality+Lying	0.361*** (0.064)	0.542*** (0.089)	0.230*** (0.067)	0.276** (0.084)	0.543*** (0.090)
Observability+Lying	0.510*** (0.058)	0.680*** (0.082)	0.330*** (0.059)	0.360*** (0.070)	0.663*** (0.084)
Constant (Externality+Choosing)	0.365*** (0.035)	0.403*** (0.050)	0.324*** (0.031)	0.399*** (0.050)	0.423*** (0.052)
Discount factor δ					
Observability+Choosing	0.046 (0.025)	0.074* (0.037)	0.066* (0.033)	0.036 (0.067)	0.017 (0.040)
Externality+Lying	0.053* (0.024)	0.088* (0.035)	0.046 (0.033)	0.071 (0.065)	0.014 (0.039)
Observability+Lying	0.046* (0.022)	0.102** (0.033)	0.062* (0.027)	0.080 (0.058)	-0.007 (0.036)
Constant (Externality+Choosing)	0.946*** (0.022)	0.884*** (0.031)	0.909*** (0.024)	0.930*** (0.051)	1.001*** (0.035)
Present bias β					
Observability+Choosing					0.138 (0.105)
Externality+Lying					0.192* (0.104)
Observability+Lying					0.315*** (0.093)
Constant (Externality+Choosing)					0.652*** (0.084)
CRRRA coefficient α					
		1.758*** (0.130)	1.206*** (0.129)	2.645*** (0.246)	1.726*** (0.128)
Stochastic error σ					
	0.453*** (0.039)	0.696*** (0.066)	0.420*** (0.052)	0.356*** (0.067)	0.691*** (0.065)
Log likelihood	-6382.03	-11785.10	-7670.77	-2064.28	-11769.6
# observations	12,825	21,015	14,517	4,266	21,015
# subjects	478	478	333	474	478

In the table above, comparisons are made against the treatment Externality+Choosing using a maximum likelihood model. Standard errors are reported in parenthesis while stars indicate significant difference from the benchmark treatments; three stars stand for $p < 0.01$, two stars for $p < 0.05$, and one star for $p < 0.09$. In the column (1) we use a utility structural form of hyperbolic discounting that does not consider risk preferences. In the column (2) we refine the estimations by incorporating risk preferences, and in particular assuming Constant Relative Risk Aversion (CRRA). In column (3) we repeat the estimation exercise but as a robustness check, we drop corner solutions, i.e., all those subjects who exhibited no switching points in any of the time choice sets. In column (4) we perform an additional robustness check in which we only consider subjects who viewed each of the three time choice sets first; that is, we only keep the choices for today's choice set from subjects who viewed this set first, the choices for 1-week's choice sets from subjects who viewed this set first, and so on. Lastly, in column (5) we assume a utility function of quasi-hyperbolic discounting, still under CRRA, so as to estimate present bias parameters.

7.3.1. Exponential Discounting with CRRA

In the table above, column (2) shows maximum likelihood estimations for exponential discounting under CRRA and while we show subjects' risk preferences in the appendix, we use this section to discuss the results of this estimation in more detail. We use the following graph to depict the results more intuitively.

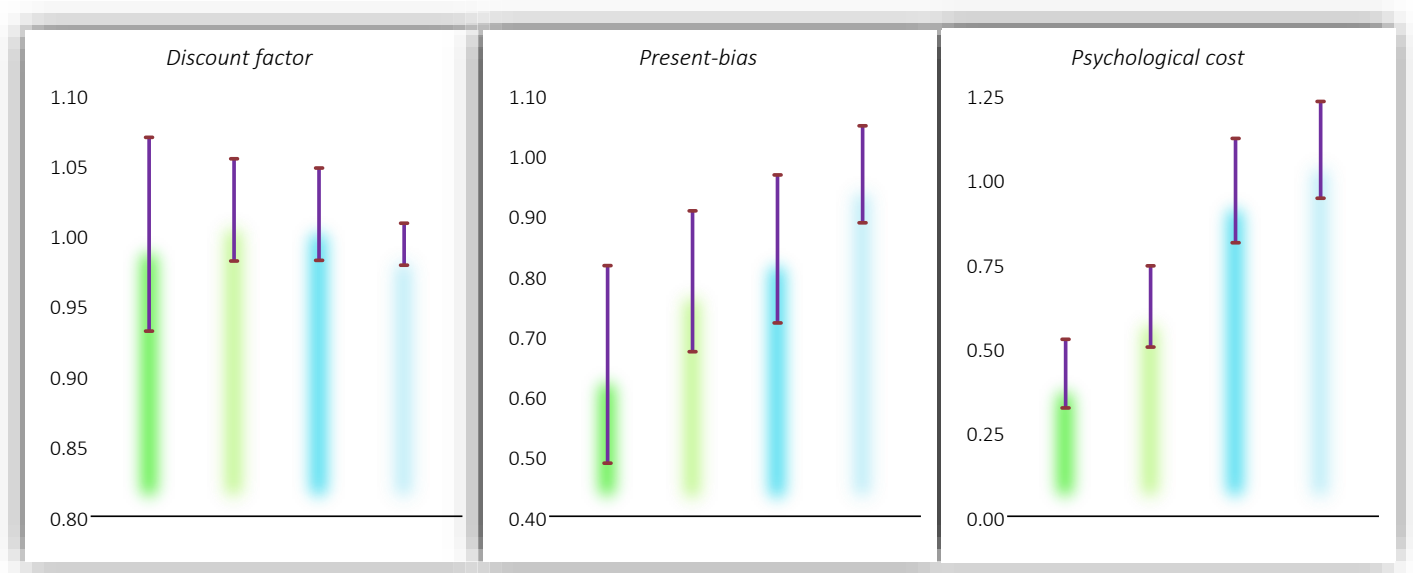


Panel B shows the estimated psychological costs for each of our four treatments. The progression of these costs is the same as we have shown in the beginning of results' section. The same state of monetary incentives produces significantly different psychological costs, starting with the lowest at the Externality+Choosing treatment, then costs increases gradually to Observability+Choosing, to Externality+Lying, and finally the costlier environment for subjects is Observability+Choosing. As shown in column (2), the differences in psychological costs between the benchmark treatment and all the other treatments is significant, as is the difference between treatment Externality+Lying and Observability+Lying.

The comparisons between treatments maintain significant differences also when we turn to compare discount factors as shown in column (2) and panel A in the graph above. However, the meaningful exercise here, is to compare the respective discount rates against 1. That is, compare if any discounting of psychological costs occurs at all (as a discounting rate indistinguishable from one suggests that future events are experienced in the same manner as present events). Such exercise returns interesting results. While future externalities and observability cause discounting of psychological costs, they do so only in the choosing conditions, not in the lying conditions. In particular, discounting factor in the Externality+Choosing is significantly different from 1 ($p < 0.01$), and so is the discounting factor in the Observability+Choosing treatment ($p = 0.018$). However discounting rates of Externality+Lying and Observability+Lying are not significantly different from 1.

7.3.2. Quasi-hyperbolic Discounting with CRRA

Turning to column (5) of the results table, we will now consider a utility function of quasi-hyperbolic discounting, meaning, we will introduce parameter beta which will show us the extent to which a time change in outcomes, from a present moment to a future one, is affecting how psychological costs are experienced. Then, as before, discount factor delta, will show us how the length of time distance is additionally affecting psychological costs. We represent graphically the results from column (5) below.



As with the hyperbolic estimations, psychological costs follow the same pattern of significant differences also with the quasi-hyperbolic estimations, as seen in panel C above. Equivalent decisions become progressively costlier when we move from Externality+Choosing treatment to Observability+Lying. However, the beta delta model we consider here projects a clearer story. It seems that delta, the length of time, does not affect costs, however, beta, the sensitivity of costs to any time shift from present time, is what driving the diverseness across treatments. While no discount rate across the four treatments is now significantly different from 1, present bias is in all but the Observability+Lying treatment⁷.

In summary, the results from this section are the following. In the treatments with choosing, observability imposes a significantly higher psychological cost. The psychological cost, in both these treatments, becomes smaller when the externalities and observability move from today to the future. However, how further in the future these outcomes move, does not affect costs in a significant manner. When we turn to the treatments with lying, we see that lying per se, adds an additional, significant, cost in both Externality and Observability conditions⁸. Further, similarly to the choosing conditions, Observability significantly increases further the psychological cost of lying⁹. These lying costs exhibit intertemporal variations. While the distance from the present time in which Externalities and Observability occur does not affect either cost, moving these consequences from the present time reduces the lying costs without

⁷ With the following p values for each treatment: Externality+Choosing $p < 0.01$, Observability+Choosing $p < 0.01$, Externality+Lying $p = 0.013$.

⁸ Comparing Externality+Choosing to Externality+Lying, and Observability+Choosing to Observability+Lying respectively.

⁹ Comparing Externality+Lying to Observability+Lying

observability. Said differently, when externalities move in the future, lying costs become smaller. However, when lying is observable, lying costs never ease, even if observability occurs in the future.

8. Conclusion

Lying and deception have always been morally loaded actions across many cultures and social systems. Their presence imposes particular problems for social wellbeing as it erodes trust in public functions and institutions. While the problems that come from lying and cheating are not new, this era, being one of fast information exchange and dense social networks, provides a particularly fertile ground for the spur of dishonest doing. Cases of fake news, or manipulative truth distortion can have many more than financial consequences. For example, during the Covid pandemic, there were many instances of individuals fabricating negative test results or vaccination records leading to worsening the spread of the virus, potentially leading to more deaths. Equally disturbing were the cases of fake news and conspiracy theories surrounding Covid vaccines which fueled the resistance of many individuals to vaccinate and protect themselves and others.

These examples serve to demonstrate that lying, while a primal practice, has not been studied thoroughly enough for society to advance successful incentive schemes in order to confine its presence. This paper discussed how pivotal the intertemporal components of lying and social preferences are and marked an important breach in the literature. As a first step towards addressing this breach, we begin by studying the intertemporal aspects of the psychological costs of lying. We constructed an experimental environment where lying opportunities arise and introduced asynchronies between lying decisions and lying externalities and observability. Aiming to assess the extent to which our findings are endemic to lying or whether they are driven by antisocial preferences in general, we extended our experimental set up to consider an equivalent environment in which lying opportunities are removed and replaced with simple distributive choices. We analyze our data using maximum likelihood estimations of utilities with varying structural forms of utility- from exponential to quasi-hyperbolic discounting models.

We report significant behavioral differences when outcomes occur as products of lying rather than distributive choices which indicates that lying is particular category of antisocial behavior, giving rise to independent psychological costs. Further, the results from the structural estimations suggest that the costs we experience when our lies cause externalities for others are rather time invariant. That is, we are not strongly inclined to lie less if the consequences from our lies occur in future timings. This is effect of

experiencing psychological costs the moment we lie rather than the moment the consequences of our lies occur, is particularly stronger in the cases where lying is observable to others. When lying is attached to observability, the psychological costs endured are not attenuated if observability is distanced from the moment lying occurs. However, incentive equivalent choices that do not involve lying bring about different intertemporal decisions as the costs from future externalities and observability are subject to significant discounting and present bias.

References

- Abeler, Johannes, Anke Becker, and Armin Falk. 2014. "Representative evidence on lying costs." *Journal of Public Economics* 113: 96-104. doi:dx.doi.org/10.1016/j.jpubeco.2014.01.005.
- Abeler, Johannes, Daniele Nosenzo, and Collin Raymond. 2019. "Preferences for truth-telling." *Econometrica* 1115-1153.
- Akerlof, George A, and Rachel E Kranton . 2000. "Economics and identity." *Quarterly journal of Economics* 715-753.
- Aksoy, Billur, and Marco Palma. 2018. "The Effects of Scarcity on Cheating and In-Group Favoritism." *No 20181111-001, Working Papers, Texas A&M University*.
<https://EconPapers.repec.org/RePEc:txm:wpaper:20181111-001>.
- Andersen, S., U. Gneezy, A. Kajackaite, and J. Marx. 2018. "Allowing for reflection time does not change behavior in Dictator and Cheating Games." *Journal of Economic Behavior & Organization* 24-33.
- Andersen, Steffen, Glenn W Harrison, Lau I. Morten, and Elisabet E Rutström. 2008. "Eliciting Risk and Time Preferences." *Econometrica*. <https://doi.org/10.1111/j.1468-0262.2008.00848.x>.
- Antinyan, A., T. Bassetti, L. Corazzini, and F. Pavesi. 2020. "Trust in the Healthcare System and COVID-19 Treatment in the Developing World. Survey and Experimental Evidence from Armenia." *No 2020:10, Working Papers, Department of Economics, Un. of Venice, "Ca'Foscari"*.
- Attanasi, Giuseppe, Fortuna Casoria, Samuele Centorrino, and Giulia Urso. 2013. "Cultural investment, local development and instantaneous social capital: A case study of a gathering festival in the South of Italy." *The Journal of Socio-Economics* 228-247.
- Banfield, Edward C., and L. Fasano. 1958. *The Moral Basis of a Backward Society*. The Free Press.
- Barr, Abigail , and Georgia Michailidou. 2017. "Complicity without Connection." *Journal of Economic Behavior & Organization* 142: 1-10. doi:<https://doi.org/10.1016/j.jebo.2017.07.013>.

- Battiston, Pietro, and Simona Gamba. 2016. *When the two ends meet: an experiment on cooperation across the Italian North-South divide*. Pisa: Laboratory of Economics and Management (LEM), Sant'Anna School of Advanced Studies,.
- Bernhard, Helen, Ernst Fehr, and Urs Fischbacher. 2006. "Group affiliation and altruistic norm enforcement." *The American Economic Review* 96 (2): 217-221.
- Bicchieri, Cristina, Eugen Dimant, and Silvia Sonderegger. 2020. "It is not a lie if you believe the norm does not apply." *available at SSRN*. <https://ssrn.com/abstract=3326146>.
- Bickel, Warren K, Mikhail N. Koffarnus, Lara Moody, and George Wilson. 2014. "The behavioral- and neuro-economic process of temporal discounting: A candidate behavioral marker of addiction." *Neuropharmacology* 518-527. <https://doi.org/10.1016/j.neuropharm.2013.06.013>.
- Bigoni, Maria, Stefania Bortolotti, and Marco Casari. 2017. *At the root of the North-South cooperation gap in Italy: Preferences or beliefs?* Quaderni - Working Paper DSE N° 1092 .
- Bigoni, Maria, Stefania Bortolotti, Marco Casari, Diego Gambetta, and Francesca Pancotto. 2016. "Amoral Familism, Social Capital, or Trust? The Behavioural Foundations of the Italian North-South Divide." *The Economic Journal* 1318-1341.
- Bowles, Samuel. 1998. "Endogenous preferences: The cultural consequences of markets and other economic institutions." *Journal of economic literature* 36 (1): 75-111.
- Butler, Jeffrey V, Danila Serra, and Giancarlo Spagnolo. 2019. "Motivating Whistleblowers." *Management Science*. doi:<https://doi.org/10.1287/mnsc.2018.3240>.
- Cadsby, C. Bram, Ninghua Du, and Song Fei . 2016. "In-group favoritism and moral decision-making." *Journal of Economic Behavior & Organization* 59-71.
- Campos-Mercade, P., A. Meier, F. Schneider, and E. Wengström. 2020. "Prosociality Predicts Health Behaviors during the COVID-19 Pandemic." *University of Zurich, Department of Economics, Working Paper No. 346, 2020, Available at SSRN: <https://ssrn.com/abstract=3604094>*.
- Cialdini, Robert B., Raymond R. Reno, and Carl A. Kallgren. 1990. "A focus theory of normative conduct: Recycling the concept of norms to reduce littering in public places." *Journal of Personality and Social Psychology* 1015-1026.
- Conrads, J. , B. Irlenbusch, R. M. Rilke, and G. Walkowitz. 2013. "Lying and team incentives." *Journal of Economic Psychology* 1-7.
- Dana, Jason, Roberto A. Weber, Jason Xi Kuang, and . 2007. "Exploiting moral wiggle room: experiments." *Economic Theory* 33: 67-80. doi:10.1007/s00199-006-0153-z.
- Demel, Simona, Abigail Barr, Luis Miller, and Paloma Ubeda. 2019. "Commitment to Political Ideology is a Luxury Only Students Can Afford: A Distributive Justice Experiment." *Journal of Experimental Political Science* 6 (1): 33-43. doi:10.1017/XPS.2018.14.
- Diekmann, Andreas, Wojtek Przepiorka, and Heiko Rauhut. 2015. "Lifting the veil of ignorance: An experiment on the contagiousness of norm violations." *Rationality and Society* 309-333.

- Ederer, Florian, and Ernst Fehr. 2007. "Deception and incentives: How dishonesty undermines effort provision." *University of Zurich Institute for Empirical Research in Economics Working Paper* 341.
- Erat, Sanjiv, and Uri Gneezy. 2012. "White lies." *Management Science* 58 (4): 723-733.
- Erkal, Nisvan, Lata Gangadharan, and Nikos Nikiforakis. 2011. "Relative Earnings and Giving in a Real-Effort Experiment." *American Economic Review* 101: 3330-48. doi: 10.1257/aer.101.7.3330.
- Falk, Armin, and James J Heckman. 2009. "Lab Experiments Are a Major Source." *Science* 326 (5952): 535-538. doi:DOI: 10.1126/science.1168244.
- Fink, S. 2020. *The Hardest Questions Doctors May Face: Who Will Be Saved? Who Won't?* New York Times. <https://www.nytimes.com/2020/03/21/us/coronavirus-medical-rationing.html>.
- Fischbacher, Urs, and Franziska Föllmi-Heusi. 2013. "Lies in disguise—an experimental study on cheating." *Journal of the European Economic Association* 11 (3): 525-547.
- Fortin, B., Lacroix, G., and M. C. Villeval. 2007. "Tax evasion and social interactions." *Journal of Public Economics* 2089-2112.
- Frakt, A. 2020. *Who Should Be Saved First? Experts Offer Ethical Guidance*. New York Times. <https://www.nytimes.com/2020/03/24/upshot/coronavirus-rationing-decisions-ethicists.html?auth=login-google>.
- Fu, Feng, Corina E Tarnita, Nicholas A Christakis, Long Wang, David G Rand, and Martin A Nowak. 2012. "Evolution of in-group favoritism." *Scientific Reports*. doi:http://dx.doi.org/10.1038/srep00460.
- Gächter, Simon, and Johnathan F Schulz. 2016. "Intrinsic honesty and the prevalence of rule violations across societies." *Nature* 496-499.
- Gächter, Simon, and Jonathan F Schulz. 2016. "Intrinsic honesty and the prevalence of rule violations across societies." *Nature* 531: 496-499.
- Garbarino, Ellen, Robert Slonim, and Marie Claire Villeval. 2018. "A method to estimate mean lying rates and their full distribution." *Journal of the Economic Science Association* 4 (2): 136-150. doi:10.1007/s40881-018-0055-4.
- Gino, Francesca, and Lamar Pierce. 2010. "Lying to level the playing field: Why people may dishonestly help or hurt others to create equity." *Journal of Business Ethics* 95 (1): 89-103.
- Gino, Francesca, Shahar Ayal, and Dan Ariely. 2013. "Self-serving altruism? The lure of unethical actions that benefit others." *Journal of economic behavior & organization* (93): 285-292.
- Gino, Francesca, Shahar Ayal, and Dan Ariely. 2009. "Contagion and Differentiation in Unethical Behavior: The Effect of One Bad Apple on the Barrel." *Psychological Science* 393-398.
- Gneezy, Uri. 2005. "Deception: The role of consequences." *The American Economic Review* 95 (1): 384-394.
- Gneezy, Uri, Agne Kajackaite, and Joel Sobel. 2018. "Lying aversion and the size of the lie." *American Economic Review* 419-53.

- Gneezy, Uri, Bettina Rockenbach, and Marta Serra-Garcia. 2013. "Measuring Lying Aversion." *Journal of Economic Behavior & Organization*.
- Guiso, Luigi, Paola Sapienza, and Luigi Zingales. 2016. "Long-term persistence." *Journal of the European Economic Association* 1401-1436.
- Hainmueller, J., D. Hangartner, and T. Yamamoto. 2015. "Validating vignette and conjoint survey experiments against real-world behavior." *PNAS* 112 (8): 2395-2400. doi: <https://doi.org/10.1073/pnas.1416587112>.
- Hey, John D, and Chris Orme. 1994. "Investigating Generalizations of Expected Utility Theory Using Experimental Data." *Econometrica*. <https://doi.org/10.2307/2951750>.
- Hildreth, John Angus , Francesca Gino, and Max Bazerman. 2016. "Blind loyalty? When group loyalty makes us see evil or engage in it." *Organizational Behavior and Human Decision Processes* 132: 16-36.
- Holt, Charles A, and Susan K Laury. 2002. "Risk Aversion and Incentive Effects." *American Economic Review*. doi:DOI: 10.1257/000282802762024700.
- Hruschka, Daniel , Charles Efferson, Ting Jiang, Ashlan Falletta-Cowden, Sveinn Sigurdsson, Rita McNamara, Madeline Sands, Shirajum Munira, Edward Slingerland, and Joseph Henrich. 2014. "Impartial Institutions, Pathogen Stress and the Expanding Social Network." *Human Nature* 567–579.
- Huang, K., J.D. Greene, and M. Bazerman. 2019. "Veil-of-ignorance reasoning favors the greater good." *Proceedings of the National Academy of Sciences* 48 (116): 23989-23995.
- Huang, K., R. Bernhard, N. Barak-Corren, M. Bazerman, and J.D. Greene. 2020. " Veil-of-Ignorance Reasoning Favors Allocating Resources to Younger Patients During the COVID-19 Crisis." Available at: <https://doi.org/10.31234/osf.io/npm4v>.
- Hurkens, Sjaak,, and Navin Kartik. 2009. "Would I lie to you? On social preferences and lying aversion." *Experimental Economics* 12 (2): 180-192. doi:10.1007/s10683-008-9208-2.
- Ichino, Andrea, and Giovanni Maggi. 2000. "Work environment and individual background: Explaining regional shirking differentials in a large Italian firm." *The Quarterly Journal of Economics* 115 (3): 1057-1090.
- J. Akesson, Jesper, S. Ashworth-Hayes, R. Hahn, R. W. Metcalfe, and I. Rasooly. 2020. "Fatalism, Beliefs, and Behaviors During the Covid-19 Pandemic." Available at SSRN: DOI 10.3386/w27245.
- Jacobsen, Catrine, Toke Reinholt Fosgaard, and David Pascual-Ezama. 2017. "Why do we lie? A practical guide to the dishonesty literature." *Journal of Economic Surveys*.
- Jiang, Ting . 2013. "Cheating in mind games: The subtlety of rules matters." *Journal of Economic Behavior & Organization* 93: 328-336.
- Jiang, Ting. 2014. "Other-regarding Preferences and Other-regarding Cheating Cheating. Experimental Evidence from China, Italy, Japan and the Netherlands, Italy, Japan and the Netherlands." *Mimeo*.

- Jin, Jian-Min, Peng Bai, Wei He, Fei Wu, Xiao-Fang Liu, De-Min Han, Shi Liu, and Jin-Kui Yang. 2020. "Gender Differences in Patients With COVID-19: Focus on Severity and Mortality." *Frontiers in Public Health*. doi:<https://doi.org/10.3389/fpubh.2020.00152>.
- Kajackaite, Agne, and Uri Gneezy. 2017. "Incentives and cheating." *Games and economic behavior* 433-444.
- Kartik, N, O Tercieux, and R Holden. 2014. "Simple mechanisms and preferences for honesty." *Games and Economic Behavior* 284-290.
- Keizer, Kees, Siegwart Lindenberg, and Linda Steg. 2008. "The spreading of disorder." *Science* 1681-1685.
- Kocher, Martin G., Simeon Schudy, and Lisa Spantig. 2016. "I Lie? We Lie! Why? Experimental Evidence on Dishonesty Shift in Groups." *Management Science* 3995-4008.
- Lauer, T., and A Untertrifaller. 2019. "Conditional Dishonesty." *mimeo*.
- Lefebvre, Mathieu, Pierre Pestieau, Arno Riedl, and Marie Claire Villeval. 2015. "Tax evasion and social information: an experiment in Belgium, France, and the Netherlands." *International tax and public finance* 401-425.
- Levine, Emma E., and Maurice E. Schweitzer. 2014. "Are liars ethical? On the tension between benevolence and honesty." *Journal of Experimental Social Psychology* 107–117.
- Lewis, A., A. Bardis, C. Flint, C. Mason, N. Smith, C. Tickle, and J. Zinser. 2012. "Drawing the line somewhere: An experimental study of moral compromise." *Journal of Economic Psychology* 718-725.
- López-Pérez, R, and E Spiegelman. 2013. "Do Economists Lie More?" *Experimental Economics* 233-247.
- Marks, J. H. 2020. "COVID-19, Pandemic Triage, and the Polymorphism of Justice." *The American Journal of Bioethics* 20 (7): 103-106.
- Mazar, Nina, On Amir, and Dan Ariely. 2008. "More ways to cheat-expanding the scope of dishonesty." *Journal of Marketing Research* 45 (6): 651-653.
- Mounk, Y. 2020. *The Extraordinary Decisions Facing Italian Doctors*. The Atlantic. <https://www.theatlantic.com/ideas/archive/2020/03/who-gets-hospital-bed/607807/>.
- Müller, S., and A. R. Holger. 2020. "Economic Preferences and Compliance in the Social Stress Test of the Corona Crisis." *CEGE Working Papers, Number 391 – April 2020, Available at SSRN: <https://ssrn.com/abstract=3575633>*.
- Nunn, Nathan, and Leonard Wantchekon. 2011. "The slave trade and the origins of mistrust in Africa." *The American Economic Review* 101 (7): 3221-3252.
- Odum, Amy L. 2013. "Delay discounting: I'm a K, you're a K." *Journal of the Experimental Analysis of Behavior*. <https://doi.org/10.1901/jeab.2011.96-423>.

- Parsons, T., and E. A. Shils . 1951. *Toward a general theory of action*. Cambridge, MA: Harvard University Press.
- Peckham, Hannah, Nina M. . de Gruijter, Charles Raine, Anna Radziszewska, Coziana Ciurtin, Lucy R. Wedderburn, Elizabeth C. Rosser, Kate Webb, and Claire T. Deakin. 2021. "Male sex identified by global COVID-19 meta-analysis as a risk factor for death and ICU admission." *nature COMMUNICATIONS*. doi:<https://doi.org/10.1038/s41467-020-19741-6>.
- Platteau, J. P. 2000. *Institutions, Social Norms, and Economic Development*. Harwood: Academic Publishers & Routledge.
- Pruckner, Gerald J, and Rupert Sausgruber. 2008. "Honesty on the Streets-A Natural Field Experiment on Newspaper Purchasing." *SSRN 1277208*.
- Robert, R., N. Kentish-Barnes, A. Boyer, A. Laurent, E. Azoulay, and J. Reignier. 2020. "Ethical dilemmas due to the Covid-19 pandemic." *Annals of Intensive Care*. doi:<https://doi.org/10.1186/s13613-020-00702-7>.
- Rosaz, Julie, and Marie Claire Villeval. 2012. "Lies and biased evaluation: A real-effort experiment." *Journal of Economic Behavior & Organization* 84 (2): 537-549.
- Rosenbaum, Stephen Mark, Stephan Billinger, and Nils Stieglitz. 2014. "Let's be honest: A review of experimental evidence of honesty and truth-telling." *Journal of Economic Psychology* 45: 181-196.
- Rotondi , Valentina, and Luca Stanca. 2015. "The effect of particularism on corruption: Theory and empirical evidence." *Journal of Economic Psychology* (51): 219-235.
- Ruffle, Bradley J., and Yossef Tobol. 2014. "Honest on Mondays: Honesty and the temporal separation between decisions and payoffs." *European Economic Review* 126-135.
- Salvemini, G. 1955. *Scritti sulla questione meridionale: 1896-1955*. Einaudi.
- Saucet, Charlotte, and Marie Claire Villeval. 2018. "Motivated memory in dictator games." *GATE WP 1804*. <https://dx.doi.org/10.2139/ssrn.3170866>.
- Seres, G., A. H. Balleyer, N. Cerutti, J. Friedrichsen, and M. Süer. n.d. "Face Mask Use and Physical Distancing before and after Mandatory Masking: Evidence from Public Waiting Lines." *Available at SSRN*. doi:<https://dx.doi.org/10.2139/ssrn.3641367>.
- Serra-Garcia, M., and N. Szech. 2020. " Demand for COVID-19 Antibody Testing and Why It Should Be Free." *CESifo Working Paper, No. 8340*.
- Shalvi, S., M. J. J. Handgraaf, and C. K. W. De Dreu. 2011. "Ethical manoeuvring: Why people avoid both major and minor lies." *British Journal of Management*.
- Shalvi, Shaul, Ori Eldar, and Yoella Bereby-Meyer. 2012. "Honesty Requires Time (and Lack of Justifications)." *Psychological Science*.
- Shao, C. 2020. "The COVID trolley dilemma." *The American Journal of Surgery* 220 (3): 545-549. doi:<https://doi.org/10.1016/j.amjsurg.2020.05.012>.

- Sobel, Joel. 2020. "Lying and deception in games." *Journal of Political Economy* 128 (3).
- Starmer, Chris. 1999. "Experiments in economics: should we trust the dismal scientists in white coats?" *Journal of Economic Methodology*. doi: DOI:.
- Tabellini, Guido. 2010. "Culture and institutions: economic development in the regions of Europe." *Journal of the European Economic Association* 8 (4): 677-716.
- Thunstrom, L., M. Ashworth, D. Finnoff, and S. Newbold. 2020. "Hesitancy Towards a COVID-19 Vaccine and Prospects for Herd Immunity." Available at SSRN: <https://ssrn.com/abstract=3593098>.
- Truog, R. D., C. Mitchell, and G. Q. Daley. 2020. "The toughest triage: Allocating ventilators in a pandemic." *The New England Journal of Medicine* 21 (328): 1973-1975. doi:doi: 10.1056/NEJMp2005689.
- V. Capraro, Valerio, and H. Barcelo. 2020. "The Effect of Messaging and Gender on Intentions to Wear a Face Covering to Slow down COVID-19 Transmission." *PsyArXiv*. doi:<https://doi.org/10.31234/osf.io/tg7vz>.
- Walkowitz, Gari. 2018. "On the Validity of Probabilistic (and Cost-Saving) Incentives in Dictator Games: A Systematic Test." doi:<https://ssrn.com/abstract=3068380>.
- Weisel, Ori, and Shaul Shalvi. 2015. "On the collaborative roots of corruption." *PNAS* 10651-10656.
- Wynia, M. K., and P. D. Sottile. 2020. "Ethical Triage Demands a Better Triage Survivability Score." *The American Journal of Bioethics* 20 (7): 75-77.